

National Aeronautics and Space Administration (NASA)
Earth Observing System Data Information System (EOSDIS)

Case Study

John Hrastar

January 2003

Table of Contents

	Page
Executive Summary.....	3
Introduction.....	4
Early Concepts.....	6
Start-up.....	8
Implementation	13
Recovery.....	18
Lessons Learned.....	22
Conclusions.....	32
References.....	33

This study is sponsored by the NASA Academy of Program and Project Leadership (APPL), Code FT NASA Headquarters, and the Systems management Office at the Goddard Space Flight Center.

EXECUTIVE SUMMARY

The Earth Observing System Data and Information System (EOSDIS), one of NASA's largest, most complex projects, was started as part of the Earth Observing System (EOS) to collect, process, distribute, and archive the large amount of data that was to be generated by the EOS program and to archive and distribute NASA Earth science data generally.

The early concepts for the EOSDIS were shaped by NASA experiences with previous data systems. A large centralized data system was thought to be necessary to assure data would be processed on time, and distributed to all involved in the interdisciplinary research of Earth science. Data could not be hoarded by the Principal Investigators as had been done in some cases previously. A large aerospace contractor was seen as the ideal developer because these contractors had experience delivering these large systems to the Department of Defense (DoD). In general there was not much experience within NASA or any other civilian agency with data systems of this magnitude and complexity.

During the start-up phase, contractor phase B studies were done and requirements were generated. However, these were done in the context of traditional procurement. A contract was signed for a large centralized EOSDIS Core System (ECS) system which was to be developed to meet a fixed end item specification. The issue of a long-term (10 years) continuing development that needed to evolve, especially considering technology changes, was recognized but never fully resolved with the work in the contract. Some saw it as necessary to put blinders on and build the EOSDIS to the specifications. Others saw it as an R&D system that must and could evolve with time.

The split between the developers and the users came early in the implementation phase and it was driven by these different views of the system. There were attempts in the early to mid 1990s to reconcile these views and change the system to accommodate everyone. This plus some poor performance caused delays and budget overruns. Simultaneously, external pressures were having an affect. EOSDIS budgets were significantly reduced to save the Agency and the Office of Earth Science money for other purposes. This external reduction in funding was dramatic, reducing EOS funding by approximately 10 Billion dollars by mid-decade. The impact on the EOSDIS was significant. Most of the time, these reductions were not accompanied by corresponding reductions in requirements. Project visibility and poor performance combined to encourage reviews by many prestigious external review teams.

In the mid to late 1990s, under changing leadership, it was recognized that basic changes would be necessary if the EOSDIS was not to collapse totally. The focus was changed to a smaller, doable system that would be developed incrementally instead of a centralized system, being developed all at once, that promised everything to everybody. The main contract requirements were cut back to a core to be developed by the contractor, and innovative ways of doing some of the science data processing by the science teams were developed. The system became more distributed. Under the revisions, the EOSDIS was completed and is delivering data to users. It does not, however, have all the capabilities

originally requested and promised but neither did it cost as much as originally budgeted. Nevertheless, it is an impressive achievement. It handles over a terabyte of data per day, and includes the management of several petabytes of data, over 1.75 million lines of code, and support of over 900 data sets. It serves more than 2 million users per year.

A number of lessons can be drawn from the rather difficult path followed by the EOSDIS during its development.

INTRODUCTION

The Earth Observing System Data and Information System (EOSDIS) is one of NASA's largest, most complex systems. It was conceived and built as the control, collection, and processing system for all of the Earth science data to be collected as part of the Earth Observing System (EOS) program and as the distribution and archiving system for the Earth Science Enterprise of NASA. It was also to include the mission command and control of the EOS flight missions. It is a major break from the preceding data systems in NASA both in size and the interoperability concept. NASA had extensive experience in developing science information systems for individual missions but a multimission system with data volumes on the order produced by EOS had never been attempted. It handles more than a terabyte of data per day and includes the management of several petabytes of data, over 1.75 million lines of code, and support of over 900 datasets. The EOSDIS serves more than 2 million users each year, many of them non-scientists. This is orders of magnitude greater than the 10,000 science users originally envisioned. It is also a system significantly different from other large, complex NASA systems such as Apollo, the Shuttle, and the Space Station in that it is not a flight system and therefore outside the traditional expertise of the Agency.

The major challenge for the EOSDIS was to achieve a balance between the mission critical operational activities and the necessary progress in science and technology. The former requires a stable robust system; the latter continued evolution, innovation, and change. Striking the balance between these is the unique feature of an R&D environment as compared with a pure operational environment that might be found in other Agencies. This is a challenge, not only for EOSDIS, but for future systems as well.

The EOS program was conceived in the 1980's as an interdisciplinary approach to study Earth. This would require many remote sensing instruments on a number of large spacecraft that would repeatedly scan the environment and develop a 15 year (or more) data set. The need for the 15 year data set was driven by the desire to include one complete solar cycle. This data set would be available to the Earth science community to help characterize Earth and its changes. The results would not only increase our understanding of Earth but would provide information to others outside the science community, e.g., policymakers, to help assure sustainable development on Earth. The EOSDIS concept was developed with EOS. Its program structure was partly a reaction to some of the problems experienced on predecessor missions and partly a recognition that

an interdisciplinary approach involving multiple science communities (atmospheres, oceans, land) would require a different kind of system than had existed previously. An integrated vision was required. It is also the NASA contribution to the Global Change Data Information System (GCDIS).

The EOSDIS was developed in the 1990's along with the EOS flight program, and it is now operating with data from many spacecraft. The development, however, was anything but smooth. In addition to the environmental changes experienced by the Agency in the 1990's, the EOSDIS experienced its own development difficulties, some of these self imposed.

The purpose of this case study on EOSDIS is to help NASA managers, engineers, and scientists understand what happened during the implementation of the EOSDIS in order to be able to apply the lessons learned to future programs and projects. The study focuses on lessons learned, both successes and failures, but it is not a detailed history. There are neither detailed traces of the Level 1 requirements nor traces of the Program Operating Plan (POP) histories included in this report. The information was gathered mainly from principals involved with the development, and from some reference material. The EOSDIS development history is a long one, from roughly the early 1980s through 1999 and included a lot of changes both internal and external. The people involved remember it in different ways so the inputs, though generally consistent, are sometimes at odds with one another. It is unlikely the events and changes that took place during the development will be repeated soon. Nevertheless, a review of the development of the EOSDIS provides a number of lessons learned for future program managers and developers.

EOSDIS Structure

The EOSDIS is part of the EOS ground system. Its structure has evolved over the years from the original concept, including the renaming of some of the components. Following is a description of the present system.

The EOSDIS is divided into six major components: data capture by the Tracking and Data Relay Satellite system (TDRSS) and EOS Polar Ground System (EPGS); the EOS Data and Operations System (EDOS); the EOSDIS Backbone Network (EBNET); the EOSDIS Core System (ECS); the Distributed Active Archive Centers (DAAC); and the Science Investigator-led Processing Systems (SIPS). The TDRSS and the EPGS capture spacecraft science data, forward it to the EDOS for processing, and provide the telemetry and command link for controlling spacecraft health and safety. The EDOS separates the data by instrument, and performs the initial processing and back-up archiving. The EBNET delivers the real-time data to and from the operations control centers, and the science data to the DAACs. The ECS provides satellite and instrument command and control, and data product generation, archival, and distribution. Product generation is done using science software provided by the Principal Investigators (PI). Using the ECS, the DAACs process the raw data into useful products, handle all user product searches, requests and orders, and distribute data and information directly to the user community, primarily via the internet. The DAACs also archive the data and information. Each

DAAC focuses on the data needs of a specific segment of the user community. Instrument teams may propose to produce their standard products operationally and deliver them to the DAACs instead of having the DAACs do this. This would be done in the SIPS. The decision to use SIPS is made on a case-by-case basis. The introduction of the SIPS was a late development in response to ECS development problems and will be discussed later in the report. Most of the lessons learned were in the ECS so this report will concentrate on that area.

EARLY CONCEPTS

Experiences on previous and contemporary projects helped to shape the ideas for EOS and the EOSDIS. In the early 1980s the observation was made by a NASA Associate Administrator (AA) that the crowded geosynchronous orbit would eventually require communications spacecraft there to become larger to accommodate more communications traffic through that location. Although not for geosynchronous orbits, this idea was extended by the AA to apply to Earth science spacecraft, i.e., large multi-instrument platforms instead of many smaller spacecraft. Prior to Space Station, NASA had underway advanced studies of space platforms. Several of the initial Space Station concepts included one or more polar platforms to meet the needs for Earth observations from this low-altitude Earth orbit. A constellation of co-orbiting and polar orbiting space platforms was included for several years in the Space Station program. Polar orbiting spacecraft are ideal for many Earth observations requiring repetitive coverage so these ideas were combined, i.e., using large polar orbiting observatories to provide Earth science data.

These concepts were coinciding with an increasing interest in studying Earth. This was motivated in part by the concern about the impact of human society on Earth and the sustainable development prospects for the planet. This type of study would require a large interdisciplinary effort to understand the land, ocean, and atmosphere systems of Earth and how they interact. Other external pressures were in play also. The National Oceanic and Atmospheric Administration (NOAA) was experiencing budget pressures on the successful Landsat program which suggested cooperation between NASA and NOAA in collecting and using Earth science data¹. There were coordinated national programs such as the Global Change Research Program (GCRP) in which the U.S. was a major participant and NASA a major contributor. There was also a general expansionist environment in NASA that included an ambitious Space Station and the space science Great Observatories. These all contributed to the establishment of EOS as a large program that would use a few large platforms with many instruments to produce large 15 year data sets of Earth. The need for a 15 year data set was driven by the desire to include one complete solar cycle. The EOSDIS that would collect, process, distribute, and archive the data from the flight system was seen as the glue holding the EOS program together.

¹ NOAA could not get money for Landsat in the late 1970s to mid 1980s because there was a government push to commercialize the system.

At this time there were also pilot data systems being developed and operated for the various Earth science disciplines, e.g., the Pilot Land Data System (PLDS), the Pilot Ocean Data System (PODS), etc. These were stovepiped for the various disciplines and did not span Earth science as a whole. They were discussing distributed data systems, however, so in this sense the PLDS and PODS, e.g., were forerunners of the EOSDIS.

The program structure of the EOSDIS was influenced by concerns about preceding and contemporary missions. Up to that time, most of the data collected on flight missions was processed, and held, by the Principal Investigators (PI), who built the instruments, before the processed data was made available to other scientists. In some cases though, especially in the Nimbus program and space science missions, there was a tendency to hoard the data, write the papers, and delay sending the data to the National Space Science Data Center (NSSDC). In part, this was a response by the science community to the difficulties they encountered in trying to submit new data products to the NSSDC and other national data centers, e.g., the cost involved in hosting a new data set. The National Science Foundation (NSF), among others, complained about this. This process was seen as unacceptable for a large interdisciplinary program. About the same time the Upper Atmosphere Research Program (UARS) was being developed which featured a Central Data Handling Facility (CDHF). The UARS data went right to the CDHF where it was processed and distributed to all, i.e., it was made available to all PIs on UARS after an agreed time. This model was seen as a better one for the EOSDIS which would be at least an order of magnitude larger, and probably more, than the CDHF. The CDHF supported only one spacecraft. Another objective was that the EOSDIS data should be publicly available to all researchers not only those selected as EOS investigators. This was thought to be most reasonably done with a large centralized system.

Another factor influencing the EOSDIS program structure was the observation by the AA that when flight programs got into trouble they would often “borrow” money from the ground system and data processing system. This would leave the data processing function short of funds. There was also a determination by the users that the ground system not be ignored until after launch as had been seen to be the case in the past. The solution to these was seen to be the separation of the flight and ground system Unique Project Numbers (UPN) so the flight program manager could not access funds meant for the ground system. In this case the EOSDIS was even put into a separate Budget Line Item (BLI), which meant that NASA would have to get Congressional approval before it could move money in or out of that line. This also had the effect of making the EOSDIS more visible to later critics and more accessible to Congressional control and influence.

The EOSDIS as envisioned early was based on size, i.e., it had to be “large”. It would handle large amounts of data, a petabyte per year, in a centralized facility and distribute it to a large community. (At one point there were 550 PIs and co-Investigators (co-Is). Despite this large number, one of the conceivers still said, “There would be more discoveries than discoverers.”) This would be necessary because of the interdisciplinary work required to study land, oceans, and atmospheres. The 1986 report of the EOS Data Panel (the “black book”) listed over 100 top-level requirements and said, “Thus, for Eos to be successful, all of the requirements and systems’ attributes delineated within this

report... must be provided.” The data would be available to all at the same time. The EOSDIS would be all things to all scientists. The science communities bought into this concept. As someone would say later, “There was broad consensus and broad naiveté.” This optimistic, expansionist attitude, similar to the one on the contemporary Space Station, pointed toward a large system to be built right away, i.e., the endpoint should be in place as soon as possible. At least one scientist did caution that the first thing that should be done was to get a system working with existing Earth science data, perhaps using some of the pilot data systems being built in the mid 1980s. The NASA AA made this a requirement and EOSDIS Version Zero (V0) was developed from 1990 to 1994. Although existing Earth science data would later be folded into the EOSDIS it did not start that way.

START UP

The early studies for the EOSDIS started at a low level in the early 1980s at Headquarters. They picked up in the mid-1980s both at Headquarters, and at the Goddard Space Flight Center (GSFC), and with a concept study at the Jet Propulsion Laboratory (JPL) in 1983. Various science study panels and working groups were formed. The EOS PIs were selected in 1988. They were to deliver their algorithms, in the form of working software, to the EOSDIS. The phase A studies (~\$3M/year) continued until the phase B studies were funded in FY 1989 and 1990². The two study contractors were TRW and Hughes Aircraft. These studies were run out of the old Mission Operations and Data Systems Directorate (MO&DSD) at Goddard, where data systems were done, not out of the Flight Projects Directorate (FPD) where projects were done. The DAACs didn’t exist in the phase B studies. Originally only one DAAC was envisioned and it was to be at GSFC. Very quickly a second DAAC was added, at the Earth Resources Observing System (EROS) Data Center (EDC), to include Landsat land data³. Eventually several others were added to specialize in various types of Earth science data. In 1990 the EOSDIS Project was started in the Flight Projects Directorate. The EOSDIS received a new start in FY1991.

Some people have questioned the use of aerospace contractors to do what is essentially a large data system. In retrospect one person speculated that a large, non-aerospace, data system contractor would have been a better choice (even, perhaps, Microsoft). This was, however, a purposeful decision. The thinking was that this was a big job, the large aerospace contractors are used to doing big jobs for the DoD including data systems, so therefore, they should be able to do this job. There was also the perception that this problem had been solved in the classified world and therefore the defense contractors would be able to do it here. (Later one scientist said, “They don’t know how to do it any better than we do.”) There was also an expectation these contractors could put the necessary systems engineering talent on the program that might not have been available elsewhere nor affordable within the civil service. In fact, one of the reasons for consolidating the system into a large central system was to be able to have a large dollar

² There were complaints by some scientists that they did not have enough input into the phase B studies despite the fact there was science participation in them.

³ The EDC is a U.S. Geological Survey facility.

value on it and therefore attract the large competent aerospace contractor. Another reason for looking toward these contractors was the lack of expertise in large data systems, especially in systems engineering, within the Agency. This was true at Headquarters, at GSFC, and across the Agency. NASA had been doing large complex flight systems for years but did not have people experienced in large data systems, i.e., more than an order of magnitude larger than previous systems.

The role of technology in the development of EOSDIS was recognized as important early but never fully resolved. There were debates on whether one should go with the new, emerging technologies or use the more mature technologies and integrate the new ones in the future. It was acknowledged that for a long term, 10 year, program the technology would not remain static. There was disagreement on the technology issues and evolutionary rate. Some saw disk storage as an issue; others that this was a non-issue because that technology would evolve fast enough. Others saw processing capability as more of a challenge because of the number of products and requirements for multiple reprocessing. The difficulty in predicting the technology capability in the future was a problem. At least part of the science community thought the Project's view was too conservative and vice versa. Eventually there was an attempt to do both, i.e., start with existing technologies but then refresh with new technologies. This issue of the evolution of the system became a major one as time went on and contributed to instability in the requirements for the system. With hindsight it was realized that there was a mismatch between the technology lifecycle and the EOSDIS build cycle. The EOSDIS releases were based on a 2 to 3 year cycle because it took that long to collect and implement the requirements for a large, centralized system. Some information systems technology, on the other hand, was turning over every 6 to 9 months. Also, some research had suggested that as the percentage of new technologies in a new program approached 25% there would be a high risk of overrun. Even when the latest technology was incorporated through commercial products, inconsistencies in the standard versions across company products caused integration problems in the EOSDIS open distributed architecture. These were not good signs for the EOSDIS.

Requirements for the EOSDIS were generated during the phase A and phase B studies. This included generation of some requirements by the study contractors in consultation with the science community. Some of these requirements were very detailed and consistent with the pending procurement. It is clear in retrospect that there was a great deal of misunderstanding on the approach, and disagreement on the philosophy of development. Many in the science community saw this as a long term evolutionary development and that their requirements would evolve as the system and technology evolved. The instruments were not yet designed so they were not ready for the detailed design requirements needed for a centralized ECS. The requirements process was seen by some scientists as disconnected between the science community and the project and the science community and the contractor. Many in the science community did not have confidence that NASA knew how to handle large, evolutionary data systems. They were aware of the lack of NASA experience in this area. They thought the requirements should be flexible and the contractor should be able to innovate because the system was supposed to evolve. The EOSDIS Project, however, saw it as necessary to have a

complete set of requirements in order to specify a system that could be built under contract. The observation by some in the EOSDIS Project was that the scientists understood the science requirements, they understood the data rates, and they understood how to do algorithms. However, they didn't understand the EOSDIS functional requirements that constituted about 95% of all the EOSDIS requirements. The small number of science requirements did drive the larger number of functional requirements. The requirements were documented but the expectations exceeded the requirements. There was never a consensus between the community and the project on the requirements.

Several factors probably contributed to the instability in requirements that plagued the project the whole time. One was the long pre-acquisition and procurement period that took place during a time of rapid technological change over three years. During a good part of the time there was a procurement blackout that precluded people from discussing what needed to be discussed regarding requirements. There was probably not enough time to work out the requirements (not just write them down) prior to acquisition. More time should have been spent listening to people, in a less pressurized environment, who understood the technology. Some competent people weren't being listened to. There was not much "vertical" discussion, only "horizontal" discussion between the experts.

This different view of requirements has been described by one person as the difference between two cultures, a "box" culture and a "cloud" culture; a box being constrained with well defined boundaries, and a cloud being amorphous. The box culture, the view of most of the management, was to view a set of fixed requirements with blinders on and build a system that met those requirements; period. The project, knew that for a traditional procurement it was necessary to have a set of requirements that must be provided to the contractor for the system to be built. The cloud culture, the view of most of the scientists, was to see the requirements as changing as more knowledge was gained. Both of the views were justified in the way the EOSDIS was conceived and started. The challenge was to separate what was to be done in the "box" and what was to be done in the "cloud". This was never satisfactorily solved early in the program and would plague the development for a long time.

Late in development, in February 1996, a little over two years from the planned launch of the first EOS spacecraft, there was a freeze on the computation and storage requirements. Some interfaces had to be baselined at this time to preclude further delays which were rife throughout the development. (See the next section on implementation.) The instrument teams still did not know their final requirements but the database had been set in concrete. Some instrumenters were angry at the freeze thinking the requirements should still continue to evolve.

Looking back one contractor observed that this pre-acquisition phase was the major problem. Before the procurement everyone, science community and government, should know what they want to buy. Because of this different view on fixed versus evolving requirements they didn't.

Following the studies there was a procurement activity that lasted a couple of years. Because of the dynamic nature of the program and changing technology, the blackout period during the procurement was probably detrimental to resolving some of the issues but it could not be avoided. During this period a new NASA Administrator came onboard. At one point he sent both bidders back to revise their proposals based on cost realism, i.e., they were thought to be unrealistically low. Hughes won the EOSDIS Core System (ECS) phase C/D contract in 1992 and the contract was signed in 1993. The value was about \$800M or almost double the original bid before the reproposal. The contract was an end-item specification contract, i.e., the end item ECS was specified by the government and expected to be delivered by the contractor. The problem was that, as discussed above, there was not an agreement by all parties, government, science community, and contractor, on what the end item was. This was true because, despite the existence of the specification the views on system evolution diverged. The Administrator also directed that there be a Total System Performance Responsibility (TSPR) clause holding the contractor fully responsible for making the system work within budget even though there was work within the DAACs outside of the contractor's control. TRW won a separate contract for the EDOS in 1994.

Shortly before contract signing cost pressures were building because of other Agency problems. An external Red Team review of the EOS program was established by the Administrator. As a consequence, all the projects in the EOS program including EOSDIS, were charged with reducing their budgets by 30% with a direction there was to be minimum impact on science. These reductions required some negotiated reductions in the ECS contract requirements and reduced critically needed flexibility. The size and significance of these reductions, which were dramatic, are discussed more in the next section on implementation.

Hughes had structured an ECS development approach that included parallel development of a number of subsystems that were to come together for integration of the EOSDIS. On a schedule this is a "waterfall" of events that culminates in the final delivery of the system. This is not a bad model for the development of a familiar system, one that has been built a number of times before. It depends on having a stable, well defined specification. However, in the view of some, this one was over compartmentalized even for a waterfall development. There was not much communication between the parallel developments even inside the contractor's team. Looking back, many people have suggested this was the wrong development model especially considering the evolution desired and rapidly changing technology. The more appropriate model was said to be a "spiral" prototyping model. In this model you take a first cut at requirements, build the hardest part, test it and use it. If it is ok then build the next part. You iterate, build and test, build and test, until you spiral in to complete the development. This model was not in vogue in NASA at the time. Instead the model for both the EOSDIS and contemporary Space Station Freedom was to do it big, right away. These discussions were going on and the technology was changing rapidly as the contract was signed based on this very structured approach.

Early in 1993 an open meeting was held at the GSFC at which Hughes briefed the science community on the EOSDIS architecture they had proposed and been selected on. This architecture was based on the requirements they had been working on from early in the study and procurement process. The design was, as stated, highly centralized. It was perceived by some users to be a dead end design. There was little feedback from the community though at this meeting. There were groups within the science community that were talking but nothing got back to the project. Better communication at this time could have surfaced some problems early.

The ECS contract with Hughes was signed in March 1993. Later that summer an open Systems Requirements Review (SRR) was conducted by Hughes. The science community was represented by the EOSDIS Panel (a.k.a. Data Panel⁴) of the EOS Investigator's Working group (IWG), but also had broader attendance at this review. Although there was little feedback from the science community at the previous meeting there was a significant negative reaction at the SRR. They saw a rigid, centralized system being developed under an end item specification contract that was not at all what they were anticipating with respect to evolution. They could not see how this design could evolve. The requirements they had submitted previously were now seen as cast in stone with perhaps only some newer hardware to be infused when it became available. In an attempt to work the issue, the project held a retreat after the SRR that included the project, Hughes, and NASA Headquarters. Hughes was told to update the architecture and decentralize the design to address some of the community concerns. This was to be approved by an external committee of senior scientists and managers by December 15. Some significant architectural changes were made and there seemed to be general agreement at the December meeting which also included some external review team participants. In retrospect it seems clear the changes were not enough to assure a smooth development because there was still not a change in the basic philosophy which envisioned everything being built at once in parallel.

At approximately the same time there was a meeting between Hughes and the new NASA Administrator⁵. The motivation for the meeting is unclear. It was reported he was told Hughes could not do the job for the money, and also that some of the scientists were complaining directly to him about Hughes. There was also a lot of budget pressure at the time, especially on the Space Station and EOS; some in Congress were critical of the EOS program. He made it clear that he expected Hughes to do the job for the money that was on the contract or it would be cancelled. Not surprisingly he cited the cost realism re-bid proposal that had increased the cost from the original proposal. There are some who were at Hughes at the time who said that this was a significant meeting. The claimed intention at Hughes was to bring in some more of the experienced talent from the classified world and make the necessary changes seen to be needed to do the EOSDIS. But after the meeting with the Administrator the decision was to hunker down, listen to his words, and go back to the NASA way, i.e., the letter of the contract. The expansionist

⁴ This Data Panel was different from the one that did the 1986 report.

⁵ It should be noted the Administrator had had some responsibility for the TRW team bidding on EOSDIS before he became NASA Administrator. He, therefore, recused himself from EOSDIS matters to some degree in consequence.

environment at NASA of a few years earlier had disappeared. This meeting, while possibly significant, should not be overly emphasized. It was probably as much a symptom of the environment as anything. In this time period EOS was being impacted by significant funding reductions.

The meeting might have had another impact. There were known weaknesses in the Hughes proposal which had been identified by the Source Evaluation Board (SEB). However, the hard line taken by the Administrator and the desire to have Hughes live with their proposal meant that these weaknesses could not be corrected before the contract started. They were built into the start of the implementation phase.

Again, at the same time, the Administrator also convened a set of his former colleagues from the classified community to look at EOSDIS. They concluded that the approach was consistent with their best practices at the time.

Whatever the motivation these two events, the SRR and the meeting with the Administrator, signaled a rough start for the work ahead.

IMPLEMENTATION

The decision was made early (1990) to expand the role of the EOSDIS to include not only the data from the EOS flight program but data from all of the other Earth science missions, such as Tropical Rainfall Measuring Mission (TRMM), Landsat, UARS, etc. In other words it would go from being the Earth Observing System DIS to the Earth Science DIS (ESDIS)⁶. This was motivated by the fundamental assumption that all Earth science data from NASA missions would be used together. There would be new, and more integrated ways to accomplish Earth system science.

Shortly after the EOSDIS new start in late 1990 (FY1991), it became clear that it was essential to deliver some improvement in data systems early and to have something to show for the money being expended. EOSDIS sponsors, Congress for example, were anxious to get something delivered. This capability, termed Version 0 (V0) included a migration of much of the earlier pilot data system work (e.g., PLDS) into the EOSDIS, and depended a lot on this earlier development work. There were many motivations for V0. It would be useful whether EOS missions flew or not, one could learn about data system development by managing these heritage datasets and making them better accessible to the community, these data sets could be managed by distributed data centers (which came to be known as DAACs), EOSDIS could evolve from this development, etc. The World Wide Web (www) came into being about one year before the completion of V0 and was to some extent incorporated in it. The V0 system was completed by August 1994; it is still operating. The V0 delivery demonstrated that concentration on a narrowly focused goal could be successful.

Even fairly early in the implementation phase (1993-94) milestones were being missed and schedule was lost. By design, a lot of data item deliveries were requested early in the

⁶ The Project name was changed accordingly.

schedule by the government. Hughes could not keep up and was delivering “shell” documents, i.e., ones with little content. This caused renegotiation on content and delivery and pushed the schedule out even more. The ECS Preliminary Design Review (PDR) was held early in 1995 without a final set of requirements, which is a prerequisite for an end item specification contract.

In 1993 the ESDIS Project was moved organizationally from the Flight Projects Directorate (FPD) to the Mission Operations and Data Systems Directorate (MO&DSD) from which the original studies had been conducted. The MO&DSD, since reorganized out of existence at Goddard, was the traditional location for ground and data system development. It was institutionally structured, not project structured. The systems they developed were to support the institution, i.e., many projects⁷. There was no project management experience in the MO&DSD so the project was run by an Associate Director, formerly of the FPD, who had prior flight project management experience. This move in a way symbolized the ambivalence GSFC management had in locating and running the project. (The project would be moved, again, into the Flight Projects Directorate (FPD) in 1997.) EOSDIS is primarily a data system for which the expertise was in the MO&DSD but it is also a project and that management expertise lies in the FPD.

The Center did not have the appropriate expertise in the project management of large data systems although it had significant capability in flight project management. One person associated with one of the competing contractors observed that the government people were more listeners than the active participants that they had expected. This was a sign of the inexperience. One consequence of this lack of experience is that there have been six project managers over a period of 12 years and most of them had not had the experience that would be considered normal in the private sector for this particular job. Despite that, all of them contributed under difficult circumstances, and the later ones were instrumental in taking steps that would eventually complete the system. This lack of experience also existed at the program level at Headquarters. Although the vision was there to conceive the EOSDIS, the appropriate experience was not there for program management of large data systems. Shortly after the EOSDIS new start limited authority was provided to recruit a new EOSDIS Program Manager. Only in late 1994 was a senior position provided for this function and by then it was too late. This situation continued even after program management was moved to the Centers in 1996. Before the program management move to the Centers in 1996 there were about 10-12 people experienced in data systems (although not necessarily program management) at Headquarters. With the program move a number of these people retired or changed jobs. Only a couple moved to the Center for the data system program management, and they were not experienced program managers.

This lack of experienced leadership was most acute at the top. By the mid 1990s all the original prime movers and leaders at Headquarters had moved on. There was a lack of a program champion at Headquarters to set major policy and decide between competing

⁷ In fact while the project was in the MO&DSD there was pressure on the Project Manager to make sure he supported the institutional developments.

agendas, one who could direct the community and the project. There was no longer a division level organization focused on data systems and issues. As a result the project had to react to competing recommendations from various reviews without the executive leadership from Headquarters to decide between them.

This lack of experience was manifested early as a lack of change leadership after the contentious SRR. It was not only evident in the inexperienced government team but also in the lack of architects and leaders in the contractor team, and the lack of a unified science community leadership that could bring it together. There was a clamor by the science community to design a system that would evolve as they saw was necessary. The problem was that a rigid performance based contract was in place so change, although possible, had to be managed and backed by additional funds for change orders. The project recognized the complaints of the users and acceded to the requested changes. This was not done under controlled conditions but informally. Changes would be suggested in open meetings and accepted. Open communication is necessary among users and developers successful development but it cannot be used to direct contract changes. Too many people besides the Contracting Officer were giving direction to the contractor. The users would talk directly to the contractors. Some scientists were making requirements change requests when they had no responsibility for implementing the requests they made⁸. Neither the government nor the contractor had an adequate configuration control process. Change requests would not be acted on and some remained open for years. The inexperienced project people controlled neither the expectations of the users nor the process for changes in the contract. On top of this the contractor apparently did not push back but instead acceded to these requests. This whole process resulted in delays and cost increases. As a result none of the principals, the project, the contractor, the users got what they wanted. The project did try to make some changes as evidenced by the architectural changes made after the SRR retreat. (There was still a complaint by some scientists that the project didn't adopt any of the good ideas.) They also facilitated work at the contractor's plant using project personnel. The changes, however, were in the context of the existing contract and architecture not in significant changes to the basic approach which would have required a change management approach. Even at this late date, with the differing expectations of the parties involved, a very experienced management team should have been able to manage the conflict, regroup, identify doable objectives, and start to move toward these objectives. This did not happen; no one took control.

The ECS contractor's actions (Hughes) contributed to the difficulties. The observation was made that Hughes removed a good proposal team after the win and hired inexperienced programmers off the street instead of providing a top-flight development team for this very complex project. This is consistent with the comment from a former Hughes employee who said they did refrain from bringing in some top-flight talent from the classified world after the Hughes meeting with the Administrator, partly out of fear of overrunning the contract value. The Hughes staffing approach was also criticized in that they staffed very quickly with inexperienced people rather than more slowly with more experienced people that they might have had a chance of retaining. Mid-level,

⁸ Some suggestions could have been made as that , i.e., just suggestions, but were interpreted as direction. If this happened the project should have controlled the process.

experienced managers were seen to be in short supply. As it was the turnover rate was very high, approaching 38% at times and even higher in the mid-management ranks⁹. The efficiency suffered dramatically. At the nadir, the production rate was 3.28 lines of code (LOC)/day/person. This compares poorly with a more normal rate of 15LOC/day/person or even higher for very well motivated or talented programmers. Another former Hughes manager acknowledged that he didn't think Hughes had stepped up to the job. After contract award, Hughes brought in a project manager who had successfully managed a large information system development for another government agency. However, he was reassigned and was replaced by a project manager whose experience was neither project management nor information systems.

The external environment also changed in the mid 1990s from that in the early 1990s both in the budget and political areas. Originally EOS was to be about a \$30B program at completion or \$17B through 2000. The EOSDIS had about \$6B of this total. Originally there was also a lot of contingency money for the EOSDIS. However, in this time frame the Agency was struggling with budget cuts and Space Station overruns. The political climate changed also. There was now less support for EOS in Congress. In the early 1990s the National Space Council had its own recommendation to cut EOS because it was "too large". One of their recommendations was to reduce the optics size to reduce cost. This would have required a lower altitude. This was not done. The combination of these factors, project difficulties and external pressure, led to a series of reviews and rebaselines for the EOS program. The flight program was changed causing it to go to smaller spacecraft and smaller launch vehicles. The EOSDIS was also impacted. The \$17B for EOS through 2000 was eventually reduced, in a series of steps, through \$11B, \$8B, to about \$7B by mid-decade. The 1995 EOS Execution Phase Project Plan (EPPP) summarizes the rebaselined EOS program at \$6.1B with the EOSDIS at \$1.9B. (The ECS contract is the single largest part of the EOSDIS budget.) As the budget was reduced the requirements were not changed accordingly. The original idea was invoked, the one that says if the money you have for refresh costs could eventually buy you more capability then you ought to be able to maintain existing capability with less money. It was still necessary to collect, process, distribute, and archive the data so these functions remained but others were impacted. First the reprocessing capacity was reduced and then user services were impacted. Significant reductions were made in contingency funds because some functions, originally planned to be part of the NASA institutional budget, e.g., EDOS and network capabilities, now had to be paid for by EOSDIS. There was concern in the project at this time that the EOSDIS would be cancelled. One report had a member of one of the prestigious external review teams threatening exactly that recommendation. This contributed to the project's desire to please the community by accepting the proposed changes.

The world was changing rapidly in other respects besides budget and politics. The World Wide Web (www) was starting to come into its own. Some technology was moving faster than anticipated while network capacity was not. Some scientists, facing the difficulties of interfacing with the EOSDIS, looked at the Web, and then increased the pressure for a

⁹ This was exacerbated by the dot com boom in the mid 1990s during which any breathing programmer could be offered a lucrative job at any time.

more decentralized EOSDIS approach. What some scientists liked about the Web was that one could use it to distribute data after it was processed by scientists with the expertise to do the processing. So the web was suggested as a way to distribute the data¹⁰. A number of alternative, more decentralized architectures were studied and proposed by the science community. They were concerned about money because they thought the ECS was “eating their lunch,” which was not entirely accurate. Not coincidentally, the proposed more distributed system would also distribute the funding to the users and put control into the hands of the science community. Even though there were many self-inflicted problems in the ECS, the external budget reductions were having a significant impact. In particular, the enormous decrease in available contingency funds and the inability to accurately cost user services undermined the ability to manage to the changing environment and its opportunities.

Later in this time frame attempts were made to get a handle on the large workforce at the ECS contractors (720 full-time-equivalents (FTE) at one point) possibly reduce it to bring the doable work in line with the decreasing budget. The impact of a reduction of a large workforce in the state was brought to the attention of a Maryland Senator by the ECS contractor. This resulted in pressure on the Agency not to make any significant changes on this workforce which only increased the overall budget pressure. Ironically, the original idea to use a large central system with a large dollar value to assure the attention of a large aerospace contractor now made it difficult to avoid the political spotlight.

Whenever a high-profile, high budget government program has trouble there are bound to be a number of external reviews. EOSDIS was no different. Numerous independent review teams, including some from the National Research Council, descended on the project. One of the major objectives of some of these was to ascertain the cost of the individual services. The idea was then it would be possible to decide how to reshape the program by determining which services could be dropped and how much could be saved. This goal was never met. One of the reasons, but certainly not the only one, might have been the consolidated architecture employed in the design of the ECS. There was never a clear delineation of the cost per function. Everything seemed to be tied up in the infrastructure of the monolithic system that could not be cut. Without insight on cost per function no one seemed to be able to come up with any means to reduce cost in a rational manner. This problem is exacerbated by the big-build approach. Thus there was never a major EOSDIS restructuring based on the recommendations of the external review teams¹¹.

¹⁰ Although the web works well in a text environment it is not obvious the sophisticated data of EOS would be as well handled. One interesting observation was that the web protocol was market driven; it came from below. It was the opposite on the DIS.

¹¹ However, there was an initiative started as a result of one review that did start to recognize the importance of decentralization and innovation. The review (NRC, LaJolla CA, 1995) pushed for a Federation of Earth Science Information Partners (ESIPS) which would conduct experiments in research and applications in data systems. The emphasis was to be on innovation in science and technology. For approximately \$12M to \$15M, 24 Cooperative Agreement Notices were initiated by various groups to start this process. This Federation is self-governing and has expanded to include the DAACs. It is a continuing operation.

By the mid-1990s the EOSDIS was behind schedule and overrunning budget. There were doubts by some that anything could be delivered. For example, in 1996, despite regular reviews, it became obvious ECS Release A, destined to support TRMM, would not be ready on time. The project cancelled Release A and used some V0 and back-up capabilities being developed at the Goddard and Langley DAACs to support TRMM. Too much had been promised to the large, diverse science community. While it is easy to get agreement among developers of a small mission, the trade space in the EOSDIS had become unmanageable. There were too many common denominators that were being sought. There was a mismatch of expectations. Many in the science community saw it as an R&D project while the government managers just wanted to get something done.

RECOVERY

There were several management changes over time at both the ESDIS Project and at Hughes. A couple of the key ones happened around the same time in late 1996 and again in 1998.

In the 1993 time frame the NASA Administrator, concerned about the EOSDIS budget, asked a colleague of his from his pre-NASA days to review the EOSDIS. This person, a government employee in the classified world, concluded that the EOSDIS was destined to fail for a couple of reasons. One, it was too ambitious for the money allocated. (By this time the EOS budget had started to come down.) Another reason was that there was too much contention built-in. In his view the scientists had too forceful a voice and if the government didn't respond Congressional support would be lost. There was destructive tension. This person left the government in October 1996 and joined Hughes as a Vice President (VP). The GSFC Deputy Director then requested he come over to run the program from the contractor side.

In October 1996 another, more experienced project manager, one with several success to his credit, took over the ESDIS Project. The realization had set in that the EOSDIS would have to be restructured and more realistic goals set. They started to rein in requirements and put more discipline into the change process. They started to identify more doable objectives; to do things in a more orderly, incremental manner. For example, "By September 1997 we will demonstrate this capability." This goal, when set, was then met. They realized that everything that was promised originally could not be delivered based on the performance so far and the ongoing budget reductions. There was an attempt to define what could be delivered. What was seen as an excessive workforce level on the ECS contract was addressed. In essence what was attempted was to get more project management control of the project at the expense of some of the original goals. The ability to achieve these goals had been eroded by previous inexperienced management and a severely eroding budget.

At the same time the new Hughes VP who was now managing the ECS contract saw the program as one in need of strong systems management. (Systems management is the combination of project management and systems engineering.) He thought that they had

only a short time to turn the program around before it would be cancelled. He saw a major deficiency in the systems engineering capability of the Hughes development team. This same observation had been made years earlier by members of the government but the problem was never corrected. He hired two previous colleagues as systems engineers, again former government employees from the classified world, shortly after he came on. These two were experienced systems engineers in large data systems. There were a couple of other key hires also. This core of new people to lead the team understood large system systems engineering¹². A good working relationship was now established between the mid-level managers in ESDIS and these systems engineers.

Although more control and project discipline was applied, more doable goals set and achieved, and a project management mentality imposed, there were still problems in the 1997-1998 period. Hughes kept trying to deliver modules but had difficulty meeting schedules. There was still concern about the perceived excessive workforce. The government/contractor relationship, although more stable and realistic, was still confrontational. Structure and limits had been imposed and progress had been made but both teams understood that necessary changes still had to be made. The bottom had been reached and the basis for a turnaround was in place.

In December 1997 Raytheon purchased Hughes. The VP who brought in the new team left shortly thereafter but the key people he brought in remained. The general observation by the government people and the Hughes people who remained was that Raytheon was more concerned about better, more responsive project management. They stepped up to the issue better. One example was that a regular dialog was started between a Raytheon VP and the GSFC Deputy Director.

About the time the government started to get control of the situation (1997-1998), the Flight Operation System (FOS) failed a major test and, therefore, uncovered a major problem. The FOS was to be the operating system for all the EOS spacecraft but was being prepared first for the AM-1 (later renamed Terra) spacecraft which was originally scheduled for a 1998 launch¹³. It turns out that in the press of business surrounding the SDPS, the FOS had not been watched as closely as it should have been. A full-court press was applied to the problem by the government and the new Raytheon team. After considering a number of options, a decision was made to drop the existing system entirely and go to a new one based on a commercial system. This was assigned to Raytheon/Denver. Having made that choice, the FOS was completed within a year with the new system. The delayed FOS did contribute to an AM-1 launch delay but once the problem was uncovered a very good recovery was made by the government/Raytheon team. The new version of the FOS was termed EMOS-EOS Mission Operations System.

¹² Ironically what was now in place was a core of key Hughes systems engineers with experience on large data systems in the classified world. This was said to be the intention of Hughes early in the contract but was reversed after the meeting between Hughes upper management and the NASA Administrator after which Hughes went back to follow the letter of the contract with a different team in order not to overrun. It was also an original goal of the conceivers when they thought a large aerospace firm would be able to bring significant systems engineering talent to the project.

¹³ The FOS was subcontracted to Lockheed Martin by Hughes.

In early 1998 a new Office of Earth Science Associate Administrator (AA) was appointed. He had a better understanding of the EOSDIS than most of his predecessors. His experience included development of the PLDS. His thinking for EOSDIS recovery was to go back to the philosophy that resulted in the successful PLDS development. This model was a concentric one. At the core would be the tools that would be needed by anyone who used the EOSDIS such as search tools, formats, retrieval systems, etc. The next outer concentric circle would be things that would be common to two or more groups of users. At the outside ring would be very specific items needed by only a few. This idea of focusing on a small doable core was the same one being worked by the ESDIS Project at the time. Later in 1998 a new ESDIS Project Manager was assigned; one with more data system experience than her predecessors. She was instrumental in carrying this core idea into implementation. Her leadership was a key factor in achieving the cooperation necessary between the EOSDIS project team and the new contractor Raytheon.

A large budget overguide submitted by GSFC in 1998 for the EOSDIS triggered some major changes by the new AA. The directions from Headquarters to the Project were now to do a design-to-cost EOSDIS. This was necessary because of the now very constrained budget resource and the need to produce an EOSDIS before all support was lost in the community and in Congress. These directions coincided with another external review and critical report by the Littles' Team. While the Littles' review was just one of many of the EOSDIS, it came at a time when the Agency was considering organizational change. Shortly after the Littles' review there was a program reorganization. This reorganization was to change the GSFC program management structure to more closely resemble the one directed by the NASA Procedure and Guideline (NPG) 7120.5 entitled NASA Program and Project Management Processes and Requirements. Thus the EOSDIS was put under a program office in the Flight Programs and Projects Directorate (FPPD) at Goddard. (Formerly the Flight Projects Directorate (FPD).)

The recognized start of the turnaround was the development of options by the ESDIS Project, specifically Option A+, in 1998. This was at the time of the large budget overguide and the direction from the new AA to build-to-cost. Option A+ was the first time it was agreed that to get to completion things would have to put in a "box" (described previously in this report). There would be a more disciplined approach to completing fewer requirements. The science requirements were limited, only 17 high-level requirements, and prioritized. They would be worked on in that priority. These 17 requirements were restatements of requirements at a high level, but were stated in a prioritized order that was vetted with the science community. They were stated in terms of functional capabilities that had to be tested and verified. Thus the ECS SDPS became an incremental development with more frequent deliveries of completed capabilities, which it was not earlier. Option A+ was to form the core of the revised ECS SDPS using the new model. It created stability for the project, it convinced Congress that something would be built, and it signaled to the science community that everything originally promised could not be done.

Option A+ included restructuring the entire ECS contract with Raytheon. Some of the data processing was removed from the contract and assigned to Science Investigator-led Processing Systems (SIPS). These science-led facilities are essentially decentralized processing facilities at the science institutions. The SIPS scientists do their own requirements, algorithms, software, and processing for the instrument data in the SIPS. They are funded separately from the ECS SDPS for those tasks. The data is received from the ECS SDPS, processed, and returned to the ECS SDPS for distribution and archiving. The SIPS represent the decentralization of data processing that some had recommended in the first place in lieu of a central facility. They represent the next layer out from the core in the concentric model. The SIPS are simpler and more austere in operation compared to the original services envisioned as being provided by the original ECS SDPS concept. This most likely is due to the self-generated requirements and the funding received. The people who generate the requirements are the same ones responsible for remaining within the funding allocated to the SIPS. This is in contrast to the original model in which the requirements put on the ECS SDPS by the scientists involved did not come out of their direct funding, i.e., they did not have to pay for them. In the latter case there is little incentive to restrict requirements. In fairness to the science community, they were originally encouraged to believe that the ECS SDPS would be everything to everybody. (They were also concerned about costs and frustrated when their requests for cost accounting got nowhere.) In addition to the transfer of some data processing to SIPS some functionality was removed from the ECS SDPS. This included a lot of user services. The net result was that the core contract work was much reduced and more manageable and there was a better relationship between the some of the requirements and the funding in the SIPS. The Raytheon Program Manager at the time realized that something would have to go to get a “doable” ECS SDPS so he agreed to the reduction. This made the ECS SDPS “doable”. The loss, however, was a marked reduction in the services originally envisioned for the users

Also at this time cooperation between the government and contractor teams improved significantly. Contact was improved; a partnership was developed. The idea that “we’re in this together so let’s see what to do about it” superceded a more confrontational approach. There was also more support from Headquarters in that the AA, who understood the nature of the EOSDIS, now also, understood and supported the tough choices that had to be made. He helped protect the project from additional outside reviews that had required a considerable effort from the project previously.

Option A+ worked. Not in the sense that all the requirements of the original EOSDIS had been met; they had not. But instead, that an acceptable match had been achieved between the necessary functions to be accomplished and the resources available; the goal of any project. The EOSDIS implemented at the EDC DAAC supported the initial operations for Landsat 7 subsequent to its April 1998 launch. Additionally, the EOSDIS, with the restructured FOS (EMOS), supported the EOS AM-1(Terra) launch in December 1999.

EOSDIS continues to support EOS and other Earth science missions, along with the SIPS, in collecting, processing, distributing, and archiving data, as well as operating the flight missions. Additionally it is working with other organizations related to the EOSDIS

in the outer concentric rings discussed previously. These include the Federation of Earth Science Information Partners (ESIP, see footnote 10), and the Regional Earth Science Application Centers (RESAC).

The ESIP federation brings together government agencies, universities, non-profit organizations, and businesses to make Earth science information available to a broader community. It started in 1997, through Cooperative Agreement Notices (CAN), and uses a very deliberate and incremental approach. Type I ESIPs concentrate on producing standard products, on a strict schedule, in a highly reliable environment. These are the DAACs. Type II ESIPs are responsible for data and information products and services in support of Earth system science that are developmental in nature, where emphasis on flexibility and creativity is key to meeting the advanced research needs. Type III ESIPs are expected to provide information and services beyond the Earth science research community. (<http://esipfed.org>)

In 1998 NASA selected nine geographically distributed, academic/government/industry consortia to form seven RESACs. Each receives a grant to apply NASA's Earth science research to well defined problems of local interest. Some of these include precision farm management, landcover/use mapping, urban sprawl, and fire hazard management. (<http://www.esad.ssc.nasa.gov/resac>)

LESSONS LEARNED

The ESDIS development from roughly the mid-1980s through the 1990s took place during a period of very dynamic change within the Agency. As a result there were many external pressures on the project during the development. However, a number of internal decisions also exacerbated the problems and led to delays, cost overruns, and animosity between some of the teams. Therefore, in retrospect, there are a lot of lessons to be drawn from this development both in program and project management that can be applied to ensure smoother development of future flight and ground systems. In retrospect many of the things that should have been done differently in an earlier stage were not obvious until later. In other words, the lesson learned was after the fact. This is the reason the lessons learned are being covered after the discussion of the full story of the development has been laid out instead of being interspersed during the discussion in the appropriate section. Not unexpectedly, many of the key lessons learned were from the start-up phase. Had some things been done differently, many of the later problems could have been avoided.

Early Concepts

1. Don't overreact, or alternatively, (to mix metaphors) fight the last war or let the pendulum swing too far in the other direction.

One of the problems of the past was data hoarding and distribution delays after launch. This problem could have been solved by a much more modest approach that would have

addressed these issues without a very large, centralized system having user services that promised everything to everybody. Although there were other motivations for the large centralized system, such as the need for interoperability of large data sets, the reaction to past problems drove the design in large measure. This was also true of the separation of the EOSDIS UPN and BLI from the rest of EOS. While it seemed helpful at first it did provide extra visibility for the EOSDIS and therefore more vulnerability. The solution was more extreme than was needed to address the original issue. The size of the ECS could have been significantly reduced and science community support for processing and reprocessing guaranteed from the start by relying on the science teams to produce their standard data products rather than their delivering algorithms to a central facility. The science community had progressed in its expectations and recognition of data sharing as essential.

Start-up

2. Know what you want to build and be able to define it.

Although there was an amorphous vision that the EOSDIS would be all things to all people, from science researchers to K-12 students, there was never an overarching concept on how to implement this. Requirements were known but they were too diverse and encompassing. There was tension between the mission critical requirements (e.g., spacecraft command and control, acquiring and processing data to an acceptable level, etc.), and the evolution of technologies and progress on scientific understanding. The latter would drive reprocessing and reanalysis of data. There was never a consensus among the stakeholders on these fixed versus evolving requirements. Therefore, when the rigid end-item spec contract was developed it was not matched against an overarching system concept. That matching should have caused one to say this contract concept doesn't match our system concept. Perhaps we should be using a smaller contract with more flexibility to develop outwards. The need to define the system doesn't require a rigid end-item spec contract to cover the whole system at once. It means defining the development process in a logical manner under the overall concept.

The following three lessons are related so are thus labeled 3a, 3b, and 3c.

3a. The acquisition strategy must be tailored to yield the desired results for any system where the user needs are difficult to articulate and subject to evolution and technology enhancements.

Start in the pre-acquisition phase with the definition of the conceptual system architecture based on broad system objectives. See lesson 2. Tailor the acquisition strategy and associated contract such that the development is evolutionary, i.e., comprised of multiple blocks/releases where each block contains multiple increments. This incremental development allows the users to get smarter with time, i.e., understand what works and what doesn't work.

Incremental development allows demonstration of functionality at each stage. In all the discussions with the principals involved at Headquarters, the GSFC, in the science community, and with the contractors, this was the most common theme, i.e., the most universal lesson learned. This means build-test; build-test; etc. Functionality is demonstrated at each step. The final step is putting the last piece in place not connecting several separately developed subsystems and hoping they work together. This allows for corrections on smaller increments and for a change in direction if necessary. It also allows for an interim operating capability. This was achieved in one respect with the ESDIS Project developed Version 0 (V0) operating system in 1994. The incremental development approach would facilitate the needed interaction between the science users and the developers. Changes are easier to make at the increments.

3b. A build-it-by-the-yard approach is desirable to maintain cost control but still leave flexibility for evolutionary changes for a large, complex system expected to evolve over an extended development period.

Build-by-the-yard implies starting system development with a small core common to all users and then incrementally adding more functions for more specialized users. If an end-item spec contract is used it should be as small as possible to maintain development flexibility. Room should be allowed outside the core, i.e., in the outer concentric circles, to accommodate the more unique and changing requirements. The development proceeds concentrically outward as more is learned in the incremental development described in Lesson 3a. Not only is the development done incrementally but decisions based on what is possible within the overall system concept and resources can be made based on progress. If the technology does not progress, or core functions are more costly than expected, or for any other good reason, some functions originally planned for the outer concentric developments might not be done at all. One must be able to react to new technology (especially information technology (IT)) but not expect to drive it alone or get in front of it. Procuring the EOSDIS was closer to procuring an IT service than a product. This service is always harder to specify concretely all at once.

This buy-it by-the-yard approach allows for evolution over time and space and is opposite to the waterfall or turnkey approach that was attempted. A turnkey or waterfall development is appropriate when the product is relatively well envisioned, e.g., an F-16 or a series of NOAA spacecraft that are copies of an original. It will even work when the specifications can be well defined for a science mission spacecraft. However, for a dramatically new development, e.g., EOSDIS, you don't know what you don't know. A waterfall development assumes you know a lot. You put the solid specifications into the front end and a couple of years later the product emerges.

The development of Option A+ represents application of this lesson¹⁴. Option A+ became the new core. It did not include everything originally promised but did represent the necessary basics that allowed the system to be built. Some of the other requirements not included in the new core were done in other ways, e.g., through SIPS.

¹⁴ More correctly, Option A+ was a priority based set of capabilities that resulted in an incremental development approach that was budget driven-design to cost. However, it did also represent a new core.

3c. Flexible options must be available for the outer concentric developments

Once the small core development is assured more innovative methods can be used to develop the outer concentric circles. This could take many different forms. For example, an integration contractor could be used, perhaps the same one that developed the core capability, instead of one large contract that is supposed to deliver everything. Several smaller developments could then be used for the parts of the system that would be integrated by the integration contractor. These smaller developments could be tried simultaneously and the unsuccessful ones discarded. The integration contractor could easily have been a Hughes or Raytheon but the others could have been small entrepreneurial companies or academic institutions very efficient at software for example. This is equivalent to a market driven approach, i.e., try some things and discard the ones that are unsuccessful.

With a small core in place the outer circle developments could benefit from a spiral development. This is one in which you build something and then try it. If it works go on to the next step. There is no guarantee these steps would all work. The approach would not have been problem free. It is similar to the entrepreneurial approach mentioned above. However, there would have been ample opportunity to learn and change direction, without any risk to the mission critical components, had this been used. This would have been seen as somewhat risky by some of the managers at the time who were used to end-item spec contracts all the way through so it would have taken strong leadership. Whatever the approach it seems clear that the acquisition strategy (as envisioned in the early stages to attract the right contractor) was not appropriate for this large, new, IT development especially while things were changing so rapidly.

The development of the SIPS does represent a positive application of this lesson. The SIPS are on the next outer development circle from the core; they are not part of the core. The separation of the SIPS work from the core work allowed Raytheon to concentrate on the much smaller core work. The SIPS involve only some scientists with more narrow interests than the whole EOSDIS. They are entrepreneurial in nature. The persons involved are responsible for both the requirements and the budget to accomplish the tasks. This approach reduced the contention on requirements and allowed them to stabilize. This approach also assured direct involvement of scientific experts in the design and development of data processing algorithms. Also, this eliminated the duplication of this expertise by the contractor. Had an alternative acquisition strategy been the basis of the project from the outset, the project implementation shortcomings could have been avoided and the value for the dollars spent could have been improved.

4. Control expectations; tell the truth about capabilities.

A project is designed to produce a certain product within a certain time and resource allocation. The time and resources are not unlimited so the requirements on the product are not unlimited. During the expansionist era in the late 1980's the EOSDIS was focusing on one large end product and expecting it to be done in one large step. It was

promising all things to all people. This would not have been sustainable even had the original budget remained in place. When the budget was reduced there was no hope of achieving the original promises. A critical responsibility of program and project management is to keep everyone on the same page as to what can be done within the resources available. This is a major job even under stable circumstances let alone under dynamic ones. Expectations were not controlled on EOSDIS until late, i.e., the mid to late 1990s, even though the budget had been dropping years earlier. Nor had they been controlled while the project itself was making decisions that led to delays and overruns. Constrained resources will always be a challenge for a project. However, even if there is not enough money to do all that is desired it is almost always possible to do good science for the resources that are available.

5. Choose the appropriate organizational structure, staff it accordingly, and stay with it.

The development of the EOSDIS is a complex but typical project type activity. It involves management of technical requirements, schedules, and financial resources. i.e., project activities. In this sense the project correctly belongs in the (now) Flight Programs and Projects Directorate (FPPD). However, starting the early studies in the MO&DSD was probably appropriate for a data system. Because of the technical expertise in the MO&DSD, a strong partnership between the project and that team would always be necessary. But the mid-project move from the FPD to the MO&DSD (before being moved back to the FPD in 1997) was a mistake and later acknowledged as such. Part of the reason for the move was the inability of the Center to find the right blend of data system talent and project management experience. This was a problem across the Agency. However, moving the project back and forth didn't help it. While the project was in the MO&DSD there was pressure on the project manager to make sure he supported other institutional developments within the MO&DSD.

Strong, experienced systems engineering and the associated processes for both the government project office and the contractor team are essential ingredients for success. The contracting officers similarly must be strong and experienced especially on a large program where attention to detail can result in magnified consequences.

6. Keep the flight operating system (FOS) tied to the flight segment.

An original motivation for consolidating the ground system including the FOS, data capture, processing, distribution, and archiving, was to protect its resources from raids by the flight segment if the flight segment ran into overrun problems, which in the past they often did. The need to be able to operate multiple missions from one control center was also a motivation. While consolidating the FOS with the remainder of the ground system seemed like a good idea at the time it caused some problems later. The FOS development is closely linked to the development of the flight segment with respect to operations, integration and test, procedures, software, commanding, etc. Although it is relatively easy to separate the data capture, processing, etc., from the flight segment, it is difficult to separate the FOS from the flight segment. The FOS should be developed by the team developing the flight segment even if it is turned over to another contractor later to be

operated, along with other similarly developed missions, in a common mission control center¹⁵. In this case not only was the FOS separate from the flight segment development, but it was tied very closely, in lock step, to the development of the SDPS, which was unnecessary. Also, tried and true FOS development methodologies were not used; it was built with too few development versions.

As a result, development problems were hidden until it failed a key test about a year before launch. A significant effort was mounted to solve the problem and the decision was made to drop the original system and go to a new one which was completed within a year. This decisive action did demonstrate that it was possible to make major positive changes and move quickly, when it was necessary, and achieve long term flight operation efficiencies through commonality across multiple missions.

Implementation

7. A strong systems engineering capability is needed for large complex system development.

Although this seems to be an obvious “motherhood” statement this did not exist despite the availability of some good systems engineering talent. Strong systems management, which is the combination of project management and systems engineering, is necessary for a successful development. The project manager should insist on it and the strong systems engineering leadership should support the project manager. The weakness in the requirements management and change control processes demonstrated this was not present. This weakness was recognized at Hughes in late 1996 and 1997 by the incoming manager. He hired some systems engineers who were experienced in large data systems engineering. At this point, along with all the other changes being made, the government and contractor systems engineers started working together, for example, putting more emphasis on interfaces with the users.

8. If the underlying processes are not in place you don't have a chance.

These processes are the basic ones needed to run a project, e.g., configuration management and others. A Hughes manager, when he came on board well into the project, observed that, “these didn't exist at Hughes”. They were weak on the government side also. Some engineering change proposals were seen to be 3 to 4 years old, never having been acted on. When change did start to happen and positive changes were proposed, the system was not effective in responding. The proper tools are needed for estimating and understanding costs. There needs to be a good trace between functions and costs.

¹⁵ A downside to having the FOS developed by the flight segment developer is that it would be difficult for these different developers to come up with compatible systems to be operated in a common control center. This is a valid concern that would have to be addressed by the appropriate specifications to these developers from the common control center developer.

9. Program, Project, and executive leadership must be aware of the environment. There was a technology paradigm shift that impacted the EOSDIS development which should have changed the development path had it been recognized.

At about the time the procurement process was playing out the World Wide Web was coming into its own. The earlier concept of a centralized system using 1980s technology was being challenged before the ink was dry on the contract. The project did adopt a technology evolution model but many scientists claimed it was way too conservative. In fact, the model they (scientists) proposed, which was more aggressive than the project model, was itself, too conservative. The technology price/performance curve was not being used in favor of the project.

Some new architectures were studied and proposed by the science community. It is not clear if these would have been successful had they been tried. However, at that critical juncture when things were starting to diverge, partly because of technology, it would have been useful to take stock to see if the project was going in the right direction. A change in direction would have been possible although it would have taken strong leadership.

Recovery

10. Strong leadership, at all levels, is critical for development of a new, complex, highly visible system.

In the development of a complex, long-term project the nature of the work will change as one moves from the initial vision, through formulation, implementation, and operations. The leader or leaders must understand the system and its objectives, have the ability to direct a team through the complex procedures of development, keep the team focused, keep expectations under control, and adjust to changes in the environment.

A lack of experienced leadership at all levels of all the teams contributed to the early development difficulties of the EOSDIS. Although the vision for the early EOSDIS came out of Headquarters, there was a lack of program management experience in this area. This was critical because the large size of the program required direction and communication across many diverse interfaces with many agencies. The GSFC also suffered a lack of experienced project managers in large data systems compared to the expertise available in flight systems. The science community was not unified and individuals in that community were experts in neither the new technology they were touting, nor in the contracting area they were proposing to change. Similar conditions applied in the contractor's organization.

The first and primary ingredient for successful systems management is a strong government leader, e.g., project manager. There were a number of times during development that, in hindsight, are recognized as points that a strong stand could have precluded further divergence in the project. For example, at the SRR it was clear there was a difference on how the contract was structured for development and how some

thought it should be structured. This major disagreement could have been faced and acknowledged; the “headbutting” could have been managed. “OK we need to change the contract dramatically,” or “OK how do we get to where we want to go, within resources, using this contract, understanding everyone will not get everything they want. ?” Granted this would not have been easy but it was possible.

When things did start turning around it was because the AA did accept the project’s recommendation to go with SIPS, did limit the products, and did protect the project. The later project managers did put limits on requirements, did establish structure within the project, and did involve the contractor as partners. The contractor did bring people who understood systems engineering, did work with the mid-level managers as partners, and went along with the restructuring that, although it removed work (which went to the SIPS) did allow an end point to be reached. The strong leadership was essential for the turnaround and successful completion of the EOSDIS.

11. Maintaining partnerships between the teams is necessary for a successful development.

Crucial to the success of a project, especially a large one, is the maintenance of the idea of partnership between the teams. Once issues arise one cannot get into the mode that “this is your problem, fix it,” but rather, “this is our problem how do we fix it together?” This is not to deny that there are contractual responsibilities that must met, but rather that the success of one project team depends on all the project teams being successful. Unfortunately there were times when the science community lost confidence in the project. There was not a “buy-in” from that community. Many times the government and contractor would view each other as enemies. When the going was very tough some of the contractor workers saw themselves as victims, constantly being harassed. It takes strong leadership to maintain partnerships under adverse conditions. A strong leader who understands the system and has confidence in himself or herself and his or her team can be tough when necessary but also knows how to compromise when necessary. This means a willingness to listen.

It also means being creative. For example, for the government project to succeed the contractor must succeed. This could mean, for example, restructuring a contract to provide a reasonable fee in the future (even if it is all provisional) to provide incentive that will strengthen the commitment to project success.

One of the factors in the turnaround was the strengthening of the government/contractor partnership. The leadership saw to improving this relationship.

12. A large, visible government program will often draw political attention which can impact development in its own right.

This is often unavoidable for some systems, e.g., the Shuttle or Space Station. It was not necessarily unavoidable for the EOSDIS. One of the original goals was to have a large centralized system to correct some past abuses, but also to attract a large defense

contractor who supposedly had the “classified world” experience to this job. In hindsight there is compelling evidence that an incrementally developed, decentralized system would have been more effective in the long run.

Contractors, especially ones with a large workforce involved, are good lobbyists. This is another reason to maintain the partnership discussed previously. It is “our problem” and if Congress has to be briefed it can be done by both the government and contractor working as a team.

13. Endless reviews do not help a project that is in trouble.

Because of all the money and people involved, many talking directly to the Administrator and Congressional people, and because of the difficulties encountered, it was natural to attract a lot of attention. One consequence of this was an extraordinary number of external reviews imposed that required project response. This added even more to an already crowded plate. These included Government Accounting Office, Inspector General, NRC, NASA, as well as normal project reviews. Between 1993 and 1997 they had averaged about one major review per month, quite often with conflicting recommendations. The ESDIS Project was, therefore, continuously either preparing for a review or answering actions from one in addition to trying to do their regular work. The reviews themselves did not cause the turnaround. This was done by determined leadership at Headquarters and in the project. That leadership was also instrumental in limiting the number of reviews so that recovery could take place.

14. Comparison of the ESDIS Project with the Jack Lee Study factors.

In 1992 a NASA study headed by Jack Lee, former Director of the Marshall Space Flight Center (MSFC), looked at factors that drive NASA program costs and technical risks. Eight factors were identified. These factors and the characterization of the ESDIS Project with respect to them are shown in the following table.

Major factors that drive NASA program costs and technical risks.	Assessment for ESDIS
1. Inadequate Phase B definition	This was probably adequate for a typical project using an end item spec. However, the issue of which type of development approach to be used was not adequately covered and decided.
2. Unrealistic dependence on unproven technology.	<p>A more accurate description for EOSDIS might be: inadequate consideration of technology improvement and technology insertion in the program planning.</p> <p>This was a problem. There was a mismatch between the EOSDIS build releases (2-3 years) and the technology lifecycle (6-9</p>

	<p>months) that was never adequately resolved. Budgets were reduced based on new technology assumptions that did not always pan out.</p> <p>Also, it was unrealistic to assume that the application of object oriented software development methodologies, which were emerging in the early to mid 90's, could be applied against the same program baseline that was bid on a central mainframe based approach.</p>
3. Annual funding instability.	This was a major problem. There were major cuts in the budget, most without a corresponding reduction in requirements.
4. Complex organizational structure, including multiple unclear interfaces.	This was a problem. The science community was very diverse and not of one voice. The program structure and the project organizational location changed many times. For quite a while the project was left without a strong Headquarters champion to decide between competing recommendations.
5. Cost estimates that are often misused.	This was a problem. The initial funding looked to be adequate with sufficient contingency. As disagreements mounted on development philosophy between the science community and the project a call went out by some in the community to distribute the money differently. There was an attempt to associate costs and functions, which should have been possible, so one could then restructure. This couldn't be resolved, primarily because of the monolithic system design of the ECS, so there was never an agreement on how costs and functions were related.
6. Scope additions due to requirements creep.	Some requirement expansion was expected because it was supposed to be an evolutionary system. Technology changes would take care of the expansion. However, until the requirements were put in a box with Option A+ there was an unrealistic attempt to do everything for everybody in one large system that was not practical.

7. Schedule slips.	This was a problem due to some of the other factors such as underestimation of code required, skill shortages, rapidly changing technology, and inadequate processes.
8. Acquisition strategy that does not promote cost containment.	The acquisition strategy was flawed in that it provided for a very rigid end item spec type development instead of a more flexible concentric development more suited to a large IT project.

CONCLUSIONS

The EOSDIS is now an operational system. It is handling large amounts of data for the Earth Science Enterprise and continues to add to its data base as new missions come on line. There are over two million users per year, far more than the 10,000 scientists who were expected to use it originally. Ninety percent of these users are not scientists. There have been reports that say that NASA data in this area is more accessible than it once was. This was one of the original goals of the EOSDIS. It is a major NASA system and will help scientists in their study of Earth for years to come.

There are still complaints that it is overly bureaucratic to use, i.e., not user friendly. It has neither all the functionality originally envisioned nor all the user services originally planned. The reprocessing capability is lower than planned which means reprocessing time is longer. These shortcomings are the results of the rough development of the EOSDIS caused by both internal development mistakes and the changing external environment. As a consequence, the EOSDIS development had to change if it was to be viable in the end. These changes impacted the final design. The “doable” EOSDIS could not deliver all that was promised originally. Had some of the changes been implemented earlier, such as focusing on a small doable core and then expanding outward to include more innovative concepts like the SIPS, some of the difficulties that were encountered could have been avoided. The externally driven funding reductions also had a major impact.

The EOSDIS was conceived and started in one NASA era (1980s) and is being completed in an entirely different one. The fact that it is being completed at all, despite the internal and external problems, is a tribute to the tenacity of those who persevered to its completion. A number of lessons can be drawn from its development, that had they been recognized sooner, could have contributed to a smoother development. Nevertheless, these lessons can be applied to ease the development of future programs and projects.

References

1. "Report of the EOS Data Panel, Volume IIa, Earth Observing System Data and Information System," NASA Technical Memorandum 87777, 1986 ("Black Book").
2. "Execution Phase Project Plan for Earth Observing System (EOS)," 170-01-01 Revision A, Goddard Space Flight Center, May 1995.
3. "EOSDIS Alternative Architecture," Project Sequoia 2000 Team, Final Report, Contract#ECS-00012, Submitted to HAIS, September 6, 1994, Revised January 23, 1995.
4. "Special Review of the GSFC Earth Science Program/Project Management", (Littles' Review), presentation by the ESDIS Project, GSFC, June 18, 1998.
5. "Special Review of GSFC Earth Science Project Management," Wayne Littles' Team, June/July 1998.
6. GSFC Response to Special Review of GSFC Earth Science Project Management, July 30, 1998.
7. "Moving Forward: Lessons Learned," excerpted from an ESDIS Project presentation to the National Research Committee, summer 2001.
8. "A Review of the U.S. Global Change Research Program and NASA's Mission to Planet Earth Earth Observing System", National Research Council, Washington D.C., 1995.